# Learning the Parameters of Determinantal Point Process Kernels

**Raja Hafiz Affandi** RAJARA@WHARTON.UPENN.EDU
University of Pennsylvania

**Emily B. Fox** EBFOX@STAT.WASHINGTON.EDU
University of Washington

**Ryan P. Adams** RPA@SEAS.HARVARD.EDU
Harvard University

**Ben Taskar** TASKAR@CS.WASHINGTON.EDU
University of Washington

arXiv:1402.4862v1 [stat.ML] 20 Feb 2014

## Abstract

Determinantal point processes (DPPs) are well-suited for modeling repulsion and have proven useful in many applications where diversity is desired. While DPPs have many appealing properties, such as efficient sampling, learning the parameters of a DPP is still considered a difficult problem due to the non-convex nature of the likelihood function. In this paper, we propose using Bayesian methods to learn the DPP kernel parameters. These methods are applicable in large-scale and continuous DPP settings even when the exact form of the eigendecomposition is unknown. We demonstrate the utility of our DPP learning methods in studying the progression of diabetic neuropathy based on spatial distribution of nerve fibers, and in studying human perception of diversity in images.

## 1. Introduction

A determinantal point process (DPP) provides a distribution over configurations of points. The defining characteristic of the DPP is that it is a repulsive point process, which makes it useful for modeling diversity. Recently, DPPs have played an increasingly important role in machine learning and statistics with applications both in the discrete setting—where they are used as a diverse subset selection method (Affandi et al., 2012; 2013b; Gillenwater et al., 2012; Kulesza & Taskar, 2010; 2011a; Snoek et al., 2013)— and in the continuous setting for generating point configurations that tend to be spread out(Affandi et al., 2013a; Zou & Adams, 2012).

Formally, given a space $\Omega \subseteq \mathbb{R}^d$, a specific point configuration $A \subseteq \Omega$, and a positive semi-definite kernel function $L : \Omega \times \Omega \to \mathbb{R}$, the probability density under a DPP with kernel $L$ is given by

$$\mathcal{P}_L(A) \propto \det(L_A) \ , \tag{1}$$

where $L_A$ is the $|A| \times |A|$ matrix with entries $L(\mathbf{x}, \mathbf{y})$ for each $\mathbf{x}, \mathbf{y} \in A$. This defines a repulsive point process since point configurations that are more spread out according to the metric defined by the kernel $L$ have higher densities. To see this, recall that the subdeterminant in Eq. (1) is proportional to the square of the volume spanned by the kernel vectors associated with the points in $A$.

Building on work of Kulesza & Taskar (2010), it is intuitive to decompose the kernel $L$ as

$$L(\mathbf{x}, \mathbf{y}) = q(\mathbf{x})k(\mathbf{x}, \mathbf{y})q(\mathbf{y}) \ , \tag{2}$$

where $q(\mathbf{x})$ can be interpreted as the quality function at point $\mathbf{x}$ and $k(\mathbf{x}, \mathbf{y})$ as the similarity kernel between points $\mathbf{x}$ and $\mathbf{y}$. The ability to bias the quality in certain locations while still maintaining diversity via the similarity kernel offers great modeling flexibility.

One of the remarkable aspects of DPPs is that they offer efficient algorithms for inference, including computing the marginal and conditional probabilities (Kulesza & Taskar, 2012), sampling (Affandi et al., 2013a;b; Hough et al., 2006; Kulesza & Taskar, 2010), and restricting to fixed-sized point configurations ($k$-DPPs)(Kulesza & Taskar, 2011a). However, an important component of DPP modeling, learning the DPP kernel parameters, is still considered a difficult, open problem. Even in the discrete $\Omega$ setting, DPP kernel learning has been conjectured to be NP-hard (Kulesza & Taskar, 2012).

Intuitively, the issue arises from the fact that in seeking to maximize the log-likelihood of Eq. (1), the numerator yields a concave log-determinant term whereas the normalizer contributes a convex term, leading to a non-convex objective. This non-convexity holds even under various simplifying assumptions on the form of $L$.

Attempts to partially learn the kernel have been studied by, for example, learning the parametric form of the quality function $q(\mathbf{x})$ for fixed similarity $k(\mathbf{x}, \mathbf{y})$ (Kulesza & Taskar, 2011b), or learning a weighting on a fixed set of kernel experts (Kulesza & Taskar, 2011a). So far, the only attempt to learn the parameters of the similarity kernel $k(\mathbf{x}, \mathbf{y})$ has used Nelder-Mead optimization (Lavancier et al., 2012), which lacks theoretical guarantees about convergence to a stationary point.

In this paper, we consider parametric forms for the quality function $q(\mathbf{x})$ and similarity kernel $k(\mathbf{x}, \mathbf{y})$ and propose Bayesian methods to learn the DPP kernel parameters $\Theta$. In addition to capturing posterior uncertainty rather than a single point estimate, these methods can be easily modified to efficiently learn large-scale and continuous DPPs where the eigenstructures are either unknown or are inefficient to compute. In contrast, gradient ascent algorithms for maximum likelihood estimation (MLE) require kernels $L$ that are differentiable with respect to $\Theta$ in the discrete $\Omega$ case. In the continuous $\Omega$ case, the eigenvalues must additionally have a known, differentiable functional form, which only occurs in limited scenarios.

In Sec. 2, we review DPPs and their fixed-sized counterpart ($k$-DPPs). We then explore likelihood maximization algorithms for learning DPP and $k$-DPP kernels. After examining the shortcomings of the MLE approach, we propose a set of techniques for Bayesian posterior inference of the kernel parameters in Sec. 3, and explore modifications to accommodate learning large-scale and continuous DPPs. In Sec. 4, we derive a set of DPP moments assuming a known kernel eigenstructure and explore using these moments as a model-checking technique. In low-dimensional settings, we can use a method of moments approach to learn the kernel parameters via numerical techniques. Finally, we test our methods on both simulated and real-world data. Specifically, in Sec. 5 we use DPP learning to study the progression of diabetic neuropathy based on spatial distribution of nerve fibers and also to study human perception of diversity of images.

## 2. Background

### 2.1. Discrete DPPs/$k$-DPPs

For a discrete base set $\Omega = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$, a DPP defined by an $N \times N$ positive semi-definite kernel matrix $L$ is a probability measure on the $2^\Omega$ possible subsets $A$ of $\Omega$:

$$\mathcal{P}_L(A) = \frac{\det(L_A)}{\det(L + I)} \ . \tag{3}$$

Here, $L_A \equiv [L_{ij}]_{\mathbf{x}_i, \mathbf{x}_j \in A}$ is the submatrix of $L$ indexed by the elements in $A$ and $I$ is the $N \times N$ identity matrix (Borodin & Rains, 2005).

In many applications, we are instead interested in the probability distribution which gives positive mass only to subsets of a fixed size, $k$. In these cases, we consider fixed-sized DPPs (or $k$-DPPs) with probability distribution on sets $A$ of cardinality $k$ given by

$$\mathcal{P}^k_L(A) = \frac{\det(L_A)}{e_k(\lambda_1, \ldots, \lambda_N)} \ , \tag{4}$$

where $\lambda_1, \ldots, \lambda_N$ are eigenvalues of $L$ and $e_k(\lambda_1, \ldots, \lambda_N)$ is the $k$th elementary symmetric polynomial (Kulesza & Taskar, 2011a). Note that $e_k(\lambda_1, \ldots, \lambda_N)$ can be efficiently computed using recursion (Kulesza & Taskar, 2012).

### 2.2. Continuous DPPs/$k$-DPPs

Consider now the case where $\Omega \subseteq \mathbb{R}^d$ is a continuous space. DPPs extend to this case naturally, with $L$ now a kernel operator instead of a matrix. Again appealing to Eq. (1), the DPP probability density for point configurations $A \subset \Omega$ is given by

$$\mathcal{P}_L(A) = \frac{\det(L_A)}{\prod_{n=1}^{\infty}(\lambda_n + 1)} \ , \tag{5}$$

where $\lambda_1, \lambda_2, \ldots$ are eigenvalues of the operator $L$.

The $k$-DPP also extends to the continuous case with

$$\mathcal{P}^k_L(A) = \frac{\det(L_A)}{e_k(\lambda_{1:\infty})} \ , \tag{6}$$

where $\lambda_{1:\infty} = (\lambda_1, \lambda_2, \ldots)$.

In contrast to the discrete case, the eigenvalues $\lambda_i$ for continuous DPP kernels are generally unknown; exceptions include a few kernels such as the exponentiated quadratic. However, Affandi et al. (2013a) showed that a low-rank approximation to $L$ can be used to recover an approximation to a finite truncation of the eigenvalues representing an important part of the eigenspectrum. This enables us to approximate the normalizing constants of both DPPs and $k$-DPPs, and will play a crucial role in our proposed methods of Sec. 3.3.

## 3. Learning Parametric DPPs

Assume that we are given a training set consisting of samples $A^1, A^2, \ldots, A^T$, and that we model these data using a DPP/$k$-DPP with parametric kernel

$$L(\mathbf{x}, \mathbf{y}; \Theta) = q(\mathbf{x}; \Theta) k(\mathbf{x}, \mathbf{y}; \Theta) q(\mathbf{y}; \Theta) \ , \tag{7}$$

with parameters $\Theta$. We denote the associated kernel matrix for a set $A^t$ by $L_{A^t}(\Theta)$ and the full kernel matrix/operator by $L(\Theta)$. Likewise, we denote the kernel eigenvalues by $\lambda_i(\Theta)$. In this section, we explore various methods for DPP/$k$-DPP learning.

### 3.1. Learning using Optimization Methods

To learn the parameters $\Theta$ of a discrete DPP model, we can maximize the log-likelihood

$$\mathcal{L}(\Theta) = \sum_{t=1}^{T} \log\det(L_{A^t}(\Theta)) - T \log\det(L(\Theta) + I) \ . \tag{8}$$

Lavancier et al. (2012) suggests that the Nelder-Mead simplex algorithm (Nelder & Mead, 1965) can be used to maximize $\mathcal{L}(\Theta)$. This method is based on evaluating the objective function at the vertices of a simplex, then iteratively shrinking the simplex towards an optimal point. While this method is convenient since it does not require explicit knowledge of derivates of $\mathcal{L}(\Theta)$, it is regarded as a heuristic search method and is known for its failure to necessarily converge to a stationary point (McKinnon, 1998).

Gradient ascent and stochastic gradient ascent provide more attractive approaches because of their theoretical guarantees, but require knowledge of the gradient of $\mathcal{L}(\Theta)$. In the discrete DPP setting, this gradient can be computed straightforwardly, and we provide examples for discrete Gaussian and polynomial kernels in the Supplementary Material. We note, however, that these methods are still susceptible to convergence to local optima due to the non-convex likelihood landscape.

The log likelihood of the $k$-DPP kernel parameter is

$$\mathcal{L}(\Theta) = \sum_{t=1}^{T} \log\det(L_{A^t}(\Theta)) - T \log \sum_{|B|=k} \det(L_B(\Theta)) \ , \tag{9}$$

which presents an addition complication due to needing a sum over $\binom{n}{k}$ terms in the gradient.

For continuous DPPs/$k$-DPPs, gradient ascent can only be used in cases where the exact eigendecomposition of the kernel operator is known with a differentiable form for the eigenvalues (see Eq. (5)). This restricts the applicability of gradient-based likelihood maximization to a limited set of scenarios, such as a DPP with Gaussian quality function and similarity kernel. Furthermore, for kernel operators with infinite rank (such as the Gaussian), an explicit truncation has to be made, resulting in an approximate gradient of $\mathcal{L}(\Theta)$. Unfortunately, such approximate gradients are not unbiased estimates of the true gradient, so the theory associated with attractive stochastic gradient based approaches does not hold.

### 3.2. Bayesian Learning for Discrete DPPs

Instead of optimizing the likelihood to get an MLE, here we propose a Bayesian approach to that samples from the posterior distribution over kernel parameters:

$$\mathcal{P}(\Theta | A^1, \ldots, A^T) \propto \mathcal{P}(\Theta) \prod_{t=1}^{T} \frac{\det(L_{A^t}(\Theta))}{\det(L(\Theta) + I)} \tag{10}$$

for the DPP and, for the $k$-DPP,

$$\mathcal{P}(\Theta | A^1, \ldots, A^T) \propto \mathcal{P}(\Theta) \prod_{t=1}^{T} \frac{\det(L_{A^t}(\Theta))}{e_k(\lambda_1(\Theta), \ldots, \lambda_N(\Theta))}. \tag{11}$$

Here, $\mathcal{P}(\Theta)$ is the prior on $\Theta$. Since neither Eq. (10) nor Eq. (11) yield a closed form posterior, we resort to approximate techniques based on Markov chain Monte Carlo (MCMC). We highlight two techniques: random-walk Metropolis-Hastings (MH) and slice sampling, although other MCMC methods can be employed without loss of generality.

In random-walk MH, we use a proposal distribution $f(\hat{\Theta}|\Theta_i)$ to generate a candidate value $\hat{\Theta}$ given the current parameters $\Theta_i$, which are then accepted or rejected with probability $\min\{r, 1\}$ where

$$r = \left( \frac{\mathcal{P}(\hat{\Theta} | A^1, \ldots, A^T)}{\mathcal{P}(\Theta_i | A^1, \ldots, A^T)} \frac{f(\Theta_i | \hat{\Theta})}{f(\hat{\Theta} | \Theta_i)} \right) \ . \tag{12}$$

The proposal distribution $f(\hat{\Theta}|\Theta_i)$ is chosen to have mean $\Theta_i$. The hyperparameters of $f(\hat{\Theta}|\Theta_i)$ tune the width of the distribution, determining the average step size. See Alg. 1 of the Supplementary Material.

While random-walk MH can provide a straightforward means of sampling from the posterior, its efficiency requires tuning the proposal distribution. Choosing an aggressive proposal can result in a high rejection rate, while choosing a conservative proposal can result in inefficient exploration of the parameter space. To avoid the need to tune the proposal distribution, we can instead use slice sampling (Neal, 2003), which performs a local search for an acceptable point while still satisfying detailed balance conditions. We first describe this method in the univariate case, following

the "linear stepping-out" approach described in Neal (2003). Given the current parameter $\Theta_i$, we first sample $y \sim \text{Uniform}[0, \mathcal{P}(\Theta_i|A^1, \ldots, A^T)]$. This defines our *slice* with all values of $\Theta$ with $\mathcal{P}(\Theta|A^1, \ldots, A^T)$ greater than $y$ included in the slice. We then define a random interval around $\Theta_i$ with width $w$ that is linearly expanded until neither endpoint is in the slice. We propose $\hat{\Theta}$ uniformly in the interval. If $\hat{\Theta}$ is in the slice, it is accepted. Otherwise, $\hat{\Theta}$ becomes the new boundary of the interval, shrinking it so as to still include the current state of the Markov chain. This procedure is repeated until a proposed $\hat{\Theta}$ is accepted. See Alg. 2 of the Supplementary Material.

There are many ways to extend this algorithm to a multidimensional setting. We consider the simplest extension proposed by Neal (2003) where we use hyperrectangles instead of intervals. A hyperrectangle region is constructed around $\Theta_i$ and the edge in each dimension is expanded or shrunk depending on whether its endpoints lie inside or outside the slice. One could alternatively consider coordinate-wise or random-direction approaches to multidimensional slice sampling.

As an illustrative example, we consider synthetic data generated from a two-dimensional discrete DPP using a kernel where

$$q(\mathbf{x}_i) = \exp\left\{-\frac{1}{2}\mathbf{x}_i^\top \Gamma^{-1}\mathbf{x}_i\right\} \tag{13}$$

$$k(\mathbf{x}_i, \mathbf{x}_j) = \exp\left\{-\frac{1}{2}(\mathbf{x}_i - \mathbf{x}_j)^\top \Sigma^{-1}(\mathbf{x}_i - \mathbf{x}_j)\right\}, \tag{14}$$

where $\Gamma = \text{diag}(0.5, 0.5)$ and $\Sigma = \text{diag}(0.1, 0.2)$. We consider $\Omega$ to be a grid of 100 points evenly spaced in a $10 \times 10$ unit square and simulate 100 samples from a DPP with kernel as above. We then condition on these simulated data and perform posterior inference of the kernel parameters using MCMC. Fig. 1 shows the sample autocorrelation function of the slowest mixing parameter, $\Sigma_{11}$, learned using random-walk MH and slice sampling. Furthermore, we ran a Gelman-Rubin test (Gelman & Rubin, 1992) on 5 chains starting from overdispersed starting positions and found that the average partial scale reduction function across the four parameters to be 1.016 for MH and 1.023 for slice sampling, indicating fast mixing of the posterior samples.

### 3.3. Bayesian Learning for Large-Scale Discrete and Continuous DPPs

In the large-scale discrete or continuous settings, evaluating the normalizers $\det(L(\Theta)+I)$ or $\prod_{n=1}^{\infty}(\lambda_n(\Theta)+1)$, respectively, can be inefficient or infeasible. Even in cases where an explicit form of the truncated eigenvalues can be computed, this will only lead to approximate MLE solutions, as discussed in Sec. 3.1.

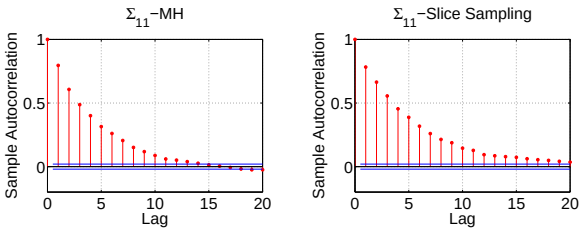


*Figure 1.* Sample autocorrelation function for posterior samples of the slowest mixing parameter of the kernel in Eq. (13) and Eq. (14) sampled using MH and slice sampling.

On the surface, it seems that most MCMC algorithms will suffer from the same problem since they require knowledge of the likelihood as well. However, we argue that for most of these algorithms, an upper and lower bound of the posterior probability is sufficient as long as we can control the accuracy of these bounds. In particular, denote the upper and lower bounds by $\mathcal{P}^+(\Theta|A^1, \ldots, A^T)$ and $\mathcal{P}^-(\Theta|A^1, \ldots, A^T)$, respectively. In the random-walk MH algorithm we can then compute the upper and lower bounds on the acceptance ratio,

$$r^+ = \left(\frac{\mathcal{P}^+(\hat{\Theta}|A^1, \ldots, A^T)}{\mathcal{P}^-(\Theta_i|A^1, \ldots, A^T)} \frac{f(\Theta_i|\hat{\Theta})}{f(\hat{\Theta}|\Theta_i)}\right) \tag{15}$$

$$r^- = \left(\frac{\mathcal{P}^-(\hat{\Theta}|A^1, \ldots, A^T)}{\mathcal{P}^+(\Theta_i|A^1, \ldots, A^T)} \frac{f(\Theta_i|\hat{\Theta})}{f(\hat{\Theta}|\Theta_i)}\right). \tag{16}$$

We can precompute the threshold $u \sim \text{Uniform}[0, 1]$, so we can still sometimes accept or reject the proposal $\hat{\Theta}$ even if these bounds have not completely converged. All that is necessary is for $u < \min\{1, r^-\}$ (immediately reject) or $u > \min\{1, r^+\}$ (immediately accept). In the case that $u \in (r^-, r^+)$, we can perform further computations to increase the accuracy of our bounds until a decision can be made. As we only sample $u$ once in the beginning, this iterative procedure yields a Markov chain with the exact target posterior as its stationary distribution; all we have done is "short-circuit" the computation once we have bounded the acceptance ratio $r$ away from $u$. We show this procedure in Alg. 3 of the Supplementary Material.

The same idea applies to slice sampling. In the first step of generating a slice, instead of sampling $y \sim \text{Uniform}[0, \mathcal{P}(\Theta_i|A^1, \ldots, A^T)]$, we use a rejection sampling scheme first propose a candidate slice as

$$\hat{y} \sim \text{Uniform}[0, \mathcal{P}^+(\Theta_i|A^1, \ldots, A^T)]. \tag{17}$$

We then decide whether $\hat{y} < \mathcal{P}^-(\Theta_i|A^1, \ldots, A^T)$, in which case we know $\hat{y} < \mathcal{P}(\Theta_i|A^1, \ldots, A^T)$ and we accept $\hat{y}$ as the slice and set $y = \hat{y}$. In the case where

$\hat{y} \in (\mathcal{P}^-(\Theta_i|A^1,\ldots,A^T), \mathcal{P}^+(\Theta_i|A^1,\ldots,A^T))$, we keep increasing the tightness of our bounds until a decision can be made. If at any point $\hat{y}$ exceeds the newly computed $\mathcal{P}^+(\Theta_i|A^1,\ldots,A^T)$, we know that $\hat{y} > \mathcal{P}(\Theta_i|A^1,\ldots,A^T)$ so we reject the proposal. In this case, we generate a new $\hat{y}$ and repeat.

Upon accepting a slice $y$, the subsequent steps for proposing a parameter $\hat{\Theta}$ proceed in a similarly modified manner. For the interval computation, the endpoints $\Theta_e$ are each examined to decide whether $y < \mathcal{P}^-(\Theta_e|A^1,\ldots,A^T)$ (endpoint is not in slice) or $y > \mathcal{P}^+(\Theta_e|A^1,\ldots,A^T)$ (endpoint is in slice). The tightness of the posterior bounds is increased until a decision can be made and the interval can be adjusted, if need be. After convergence of the interval, $\hat{\Theta}$ is generated uniformly over the interval and is likewise tested for acceptance. We illustrate this procedure in Fig. 1 of the Supplementary Material.

The lower and upper bounds of the posterior probability can in fact be incorporated in many MCMC-type algorithms. This provides a convenient and efficient way to garner posterior samples assuming that tightening the bounds can be done efficiently. In our case, the upper and lower bounds for the posterior probability depends on the truncation of the kernel eigenvalues and can be arbitrarily tightened by including more terms in the truncation. In the discrete DPP/$k$-DPP settings, the eigenvalues can be efficiently computed to a specified point using methods such as power law iterations. The corresponding bounds for a $3600 \times 3600$ Gaussian kernel example are shown in Fig. 2. In the continuous setting, explicit truncation can be done when the kernel has Gaussian quality and similarity, as we show in Sec. 5.1. For other continuous DPP kernels, low-rank approximations can be used (Affandi et al., 2013a) resulting in approximate posterior samples. In contrast, a gradient ascent algorithm for MLE is not even feasible: we do not know the form of the approximated eigenvalues, so we cannot take their derivative.

Explicit forms for the posterior probability bounds of $\Theta$ for DPPs and $k$-DPPs as a function of the eigenvalue truncations follow from Prop. C.1 and C.2 combined with Eqs. (10) and (11), respectively. Proofs are in the Supplementary Material.

**Proposition 3.1** *Let $\lambda_{1:\infty}$ be the eigenvalues of kernel $L$. Then*

$$\prod_{n=1}^{M}(1+\lambda_n) \leq \prod_{n=1}^{\infty}(1+\lambda_n) \tag{18}$$

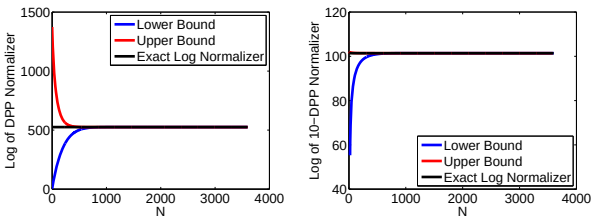


*Figure 2.* Normalizer bounds for a discrete DPP (*left*) and a 10-DPP (*right*) with Gaussian quality and similarity as in Eqs. (13) and (14) and $\Omega$ a grid of 3600 points.

*and*

$$\prod_{n=1}^{\infty}(1+\lambda_n) \leq \exp\left\{\mathrm{tr}(L) - \sum_{n=1}^{M}\lambda_n\right\}\left[\prod_{n=1}^{M}(1+\lambda_n)\right] . \tag{19}$$

**Proposition 3.2** *Let $\lambda_{1:\infty}$ be the eigenvalues of kernel $L$. Then*

$$e_k(\lambda_{1:M}) \leq e_k(\lambda_{1:\infty}) \tag{20}$$

*and*

$$e_k(\lambda_{1:\infty}) \leq \sum_{j=0}^{k} \frac{(\mathrm{tr}(L) - \sum_{n=1}^{M}\lambda_n)^j}{j!} e_{k-j}(\lambda_{1:M}) . \tag{21}$$

Finally note that the expression $\mathrm{tr}(L)$ in the bounds can be easily computed as either $\sum_{n=1}^{N} L_{ii}$ in the discrete case or $\int_\Omega L(\mathbf{x},\mathbf{x})d\mathbf{x}$ in the continuous case.

## 4. Method of Moments

Convergence and mixing of MCMC samplers can be challenging to assess. Although generic techniques such as Gelman-Rubin diagnostics (Gelman & Rubin, 1992) are applicable, we additionally provide a set of tools more directly tailored to the DPP by deriving a set of theoretical moments. When performing posterior inference of kernel parameters, we can check whether the moments of our data match the theoretical moments given by the posterior samples. This can be done in cases where the eigenstructure is fully known.

In the discrete case, we first need to compute the marginal probabilities. Borodin (2009) shows that the marginal kernel, $K$, can be computed directly from $L$:

$$K = L(I+L)^{-1} . \tag{22}$$

The $m$th moment can then be calculated via

$$\mathbb{E}[\mathbf{x}^m] = \sum_{i=1}^{N} \mathbf{x}_i^m K(\mathbf{x}_i,\mathbf{x}_i) . \tag{23}$$

In the continuous case, given the eigendecomposition of the kernel operator, $L(\mathbf{x},\mathbf{y}) = \sum_{n=1}^{\infty}\lambda_n\phi_n(\mathbf{x})^*\phi_n(\mathbf{y})$

(where $\phi_n(\mathbf{x})^*$ denotes the complex conjugate of the $n$th eigenfunction), the $m$th moment is

$$\mathbb{E}[\mathbf{x}^m] = \int_\Omega \sum_{n=1}^{\infty} \frac{\lambda_n}{\lambda_n + 1} \mathbf{x}^m \phi_n(\mathbf{x})^2 d\mathbf{x} \,. \tag{24}$$

Note that this generally cannot be evaluated in closed form since the eigendecompositions of most kernel operators are not known. However, in certain cases where the eigenfunctions are known analytically, the moments can be directly computed. For a kernel defined by Gaussian quality and similarity (see Sec. 5.1), the eigendecomposition can be performed using Hermite polynomials. In the Supplementary Material, we derive the $m$th moment in this setting.

Unfortunately, the method of moments can be challenging to use for direct parameter learning since Eqs. (23) and (39) are not analytically available in most cases. In low dimensions, these quantities can be estimated numerically, but it remains an open question as to how these moments should be estimated for large-scale problems.

## 5. Experiments

### 5.1. Simulations

We provide an explicit example of Bayesian learning for a continuous DPP with the kernel defined by

$$q(\mathbf{x}) = \sqrt{\alpha} \prod_{d=1}^{D} \frac{1}{\sqrt{\pi \rho_d}} \exp\left\{ -\frac{x_d^2}{2\rho_d} \right\} \tag{25}$$

$$k(\mathbf{x}, \mathbf{y}) = \prod_{d=1}^{D} \exp\left\{ -\frac{(x_d - y_d)^2}{2\sigma_d} \right\}, \mathbf{x}, \mathbf{y} \in \mathbb{R}^D. \tag{26}$$

Here, $\Theta = \{\alpha, \rho_d, \sigma_d\}$ and the eigenvalues of the operator $L(\Theta)$ are given by (Fasshauer & McCourt, 2012),

$$\lambda_{\mathbf{m}}(\Theta) = \alpha \prod_{d=1}^{D} \sqrt{\frac{1}{\frac{\beta_d^2+1}{2} + \frac{1}{2\gamma_d}}} \left( \frac{1}{\gamma_d(\beta_d^2+1)+1} \right)^{m_d - 1}, \tag{27}$$

where $\gamma_d = \frac{\sigma_d}{\rho_d}$, $\beta_d = (1 + \frac{2}{\gamma_d})^{\frac{1}{4}}$, and $\mathbf{m} = (m_1, \ldots, m_D)$ is a multi-index.

Furthermore, the trace of $L(\Theta)$ can be easily computed as

$$\mathrm{tr}(L(\Theta)) = \int_{\mathbb{R}^d} \alpha \prod_{d=1}^{D} \frac{1}{\pi \rho_d} \exp\left\{ -\frac{x_d^2}{2\rho_d} \right\} d\mathbf{x} = \alpha \,. \tag{28}$$

We test our Bayesian learning algorithms on simulated data generated from a 2-dimensional isotropic kernel ($\sigma_d = \sigma$, $\rho_d = \rho$ for $d = 1, 2$) using Gibbs sampling (Affandi et al., 2013a). We then learn the parameters under weakly informative inverse gamma priors on $\sigma$, $\rho$ and $\alpha$. Details are in the Supplementary Material. We tweak the $(\alpha, \rho, \sigma)$ used for simulation so that we have the following three scenarios:

(i) 10 DPP samples with average number of points=18 using $(\alpha, \rho, \sigma) = (1000, 1, 1)$

(ii) 1000 DPP samples with average number of points=18 using $(\alpha, \rho, \sigma) = (1000, 1, 1)$

(iii) 10 DPP samples with average number of points=77 using $(\alpha, \rho, \sigma) = (100, 0.7, 0.05)$.

Fig. 3 shows trace plots of the posterior samples for all three scenarios. In the first scenario, the parameter estimates vary wildly whereas in the other two scenarios, the posterior estimates are more stable. In all the cases, the zeroth and second moment estimated from the posterior samples are in the neighborhood of the corresponding empirical moments. This leads us to believe that the posterior is broad in cases where we have both a small number of samples and few points in each sample. The posterior becomes more peaked as the total number of points increases. Note that using a stationary similarity kernel allows us to garner information either from few sets with many points or many sets of few points.

**Dispersion Measure** In many applications, we are interested in quantifying the overdispersion of point process data. In spatial statistics, one standard quantity used to measure dispersion is the Ripley $K$-function (Ripley, 1977). Here, instead, we would like to use the learned parameters of the DPP to measure overdispersion as repulsion. An important characteristic of a measure of repulsion is that it should be invariant to scaling. In the Supplementary Material, we derive results that, as the data are scaled from $\mathbf{x}$ to $\eta\mathbf{x}$, the parameters scale from $(\alpha, \sigma_i, \rho_i)$ to $(\alpha, \eta\sigma_i, \eta\rho_i)$. This suggests that an appropriate scale-invariant measure of repulsion is $\gamma_i = \sigma_i / \rho_i$.

### 5.2. Applications

#### 5.2.1. Diabetic Neuropathy

Recent breakthroughs in skin tissue imaging have spurred interest in studying the spatial patterns of nerve fibers in diabetic patients. It has been observed that these nerve fibers appear to become more clustered as diabetes progresses. Waller et al. (2011) previously analyzed this phenomena based on 6 thigh nerve fiber samples. These samples were collected from 5 diabetic patients at different stages of diabetic neuropathy and one healthy subject. On average, there are 79 points in each sample (see Fig. 4). Waller et al. (2011) analyzed the Ripley $K$-function and concluded that the

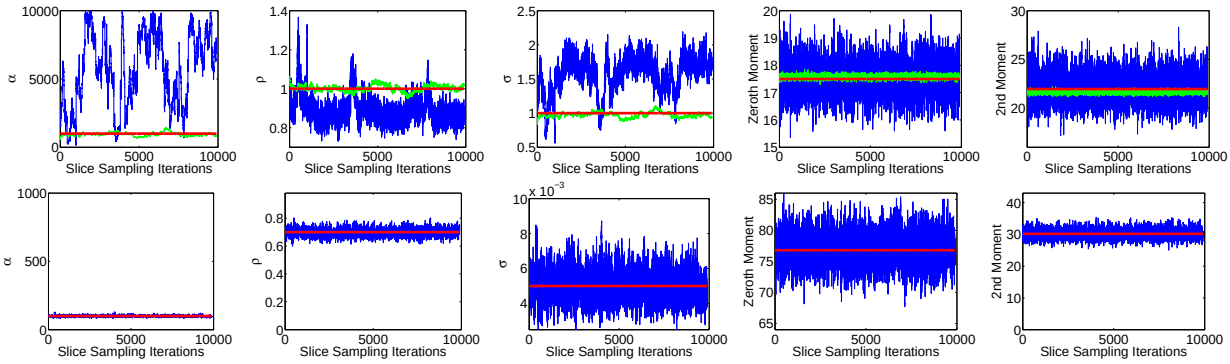


*Figure 3.* Posterior samples and moment estimates under the parameters learned for a continuous DPP with Gaussian quality and similarity. The first 3 columns show posterior samples for (from right to left) $\alpha$, $\rho$ and $\sigma$. The last two columns show the zeroth and second moment estimates. The top row are samples from Scenario (i) (blue) and Scenario (ii) (green) while the second row are samples from Scenario (iii), plotted with the same relative scales to the other scenarios. Red lines indicate the true parameter values that generated the data and their associated theoretical moments.

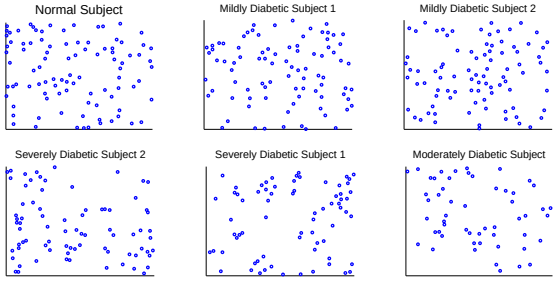


*Figure 4.* Nerve fiber samples. Clockwise: (i) Normal subject, (ii) Mildly Diabetic Subject 1, (iii) Mildly Diabetic Subject 2,(iv) Moderately Diabetic subject, (v) Severely Diabetic Subject 1 and (vi) Severely Diabetic Subject 2.

difference between the healthy and severely diabetic samples is highly significant.

We instead study the differences between these samples by learning the kernel parameters of a DPP and quantifying the level of repulsion of the point process. Due to the small sample size, we consider a 2-class study of Normal/Mildly Diabetic versus Moderately/Severely Diabetic. We perform two analyses. In the first, we directly quantify the level of repulsion based on our scale-invariant statistic, $\gamma = \sigma/\rho$ (see Sec. 5.1). In the second, we perform a leave-one-out classification by training the parameters on the two classes with one sample left out. We then evaluate the likelihood of the held-out sample under the two learned classes. We repeat this for all six samples.

We model our data using a 2-dimensional continuous DPP with Gaussian quality and similarity as in Eqs. (40) and (41). Since there is no observed preferred direction in the data, we use an isotropic kernel ($\sigma_d = \sigma$ and $\rho_d = \rho$ for $d = 1, 2$). We place weakly informative inverse gamma priors on $(\alpha, \rho, \sigma)$, as specified in the Supplementary Material, and learn the parameters using slice sampling with eigenvalue bounds as outlined in Sec. 3.3. The results shown in Fig. 5 indicate that our $\gamma$ measure clearly separates the two classes, concurring with the results of Waller et al. (2011). Furthermore, we are able to correctly classify all six samples. While the results are preliminary, being based on only six observations, they show promise for this task.

#### 5.2.2. Diversity in Images

We also examine DPP learning for quantifying how visual features relate to human perception of diversity in different image categories. This is useful in applications such as image search, where it is desirable to present users with a set of images that are not only relevant to the query, but diverse as well.

Building on work by Kulesza & Taskar (2011a), three image categories—cars, dogs and cities—were studied. Within each category, 8-12 subcategories (such as *Ford* for cars, *London* for cities and *poodle* for dogs) were queried from Google Image Search and the top 64 results were retrieved. For a subcategory subcat, these images form our base set $\Omega_{\text{subcat}}$. To assess human perception of diversity, human annotated sets of size six were generated from these base sets. However, it is challenging to ask a human to coherently select six diverse images from a set of 64 total. Instead, Kulesza & Taskar (2011a) generated a *partial result set* of five images from a 5-DPP on each $\Omega_{\text{subcat}}$ with a kernel based on the SIFT256 features (see Supplementary Material). Human annotators (via Amazon Mechanical Turk) were then presented with two images selected at random from the remaining subcategory images and asked to add the image they felt was least similar to the partial result set. These experiments resulted in about 500 samples spread evenly spread evenly across the different subcategories.

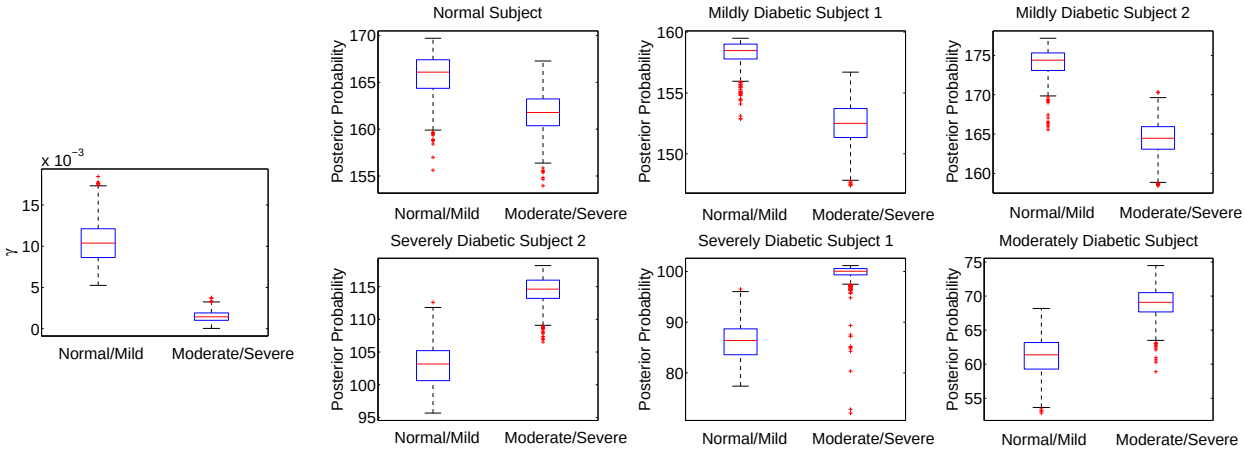


*Figure 5.* Leftmost plot shows the repulsion measure, $\gamma$, and the rest show the leave-one out log-likelihood of each sample under the two learned DPP classes: Normal/Mildly Diabetic (left box) and Moderately/Severely Diabetic (right box). The ordering of log-likelihood plots aligns with that of Fig. 4.

We aim to study how the human annotated sets differ from the top six Google results, Top-6. As in Kulesza & Taskar (2011a), we extracted three types of features from the images—color features, SIFT descriptors (Lowe, 1999; Vedaldi & Fulkerson, 2010) and GIST descriptors (Oliva & Torralba, 2006) described in the Supplementary Material. We denote these features for image $i$ as $f_i^{\mathsf{color}}$, $f_i^{\mathsf{SIFT}}$, and $f_i^{\mathsf{GIST}}$, respectively. For each subcategory, we model our data as a discrete 6-DPP on $\Omega_{\mathsf{subcat}}$ with kernel

$$L_{i,j}^{\mathsf{subcat}} = \exp\left\{-\sum_{\mathsf{feat}} \frac{\|f_i^{\mathsf{feat}} - f_j^{\mathsf{feat}}\|_2^2}{\sigma_{\mathsf{feat}}^{\mathsf{cat}}}\right\} \tag{29}$$

for $\mathsf{feat} \in \{\mathsf{color}, \mathsf{SIFT}, \mathsf{GIST}\}$ and $i, j$ indexing the 64 images in $\Omega_{\mathsf{subcat}}$. Here, we assume that each category has the same parameters across subcategories, namely, $\sigma_{\mathsf{feat}}^{\mathsf{cat}}$ for $\mathsf{subcat} \in \mathsf{cat}$ and $\mathsf{cat} \in \{\mathsf{cars}, \mathsf{dogs}, \mathsf{cities}\}$.

To learn from the Top-6 images, we consider the samples as being generated from a 6-DPP. To emphasize the human component of the 5-DPP + human annotation sets, we examine a conditional 6-DPP (Kulesza & Taskar, 2012) that fixes the five images from the partial results set and only considers the probability of adding the human annotated image. The Supplementary Material provides details on this conditional $k$-DPP.

All subcategory samples within a category are assumed to be independent draws from a DPP defined on $\Omega_{\mathsf{subcat}}$ with kernel $L^{\mathsf{subcat}}$ parameterized by a shared set of $\sigma_{\mathsf{feat}}^{\mathsf{cat}}$, for $\mathsf{subcat} \in \mathsf{cat}$. As such, each of these samples equally informs the posterior of $\sigma_{\mathsf{feat}}^{\mathsf{cat}}$. We perform posterior sampling of the 6-DPP or conditional 6-DPP kernel parameters using slice sampling with weakly informative inverse gamma priors on the $\sigma_{\mathsf{feat}}^{\mathsf{cat}}$. Details are in the Supplementary Material.

Fig. 6 shows a comparison between $\sigma_{\mathsf{feat}}^{\mathsf{cat}}$ learned from the human annotated samples (conditioning on the 5-DPP partial result sets) and the Top-6 samples for different categories. The results indicate that the 5-DPP + human annotated samples differs significantly from the Top-6 samples in the features judged by human to be important in diversity in each category. For cars and dogs, human annotators deem color to be a more important feature for diversity than the Google search engine based on their Top-6 results. For cities, on the other hand, the SIFT features are deemed important for diversity by human annotators, while the Google search engine puts a much lower weight on them. Keep in mind, though, that this result only highlights the diversity components of the results while ignoring quality. In real life applications, it is desirable to combine both the quality of each image (as a measure of relevance of the image to the query) and the diversity between the top results. Regardless, we have shown that DPP kernel learning can be informative of judgements of diversity, and this information could be used (for example) to tune search engines to provide results more in accordance with human judgement.

## 6. Conclusion

Determinantal point processes have become increasingly popular in machine learning and statistics. While many important DPP computations are efficient, learning the parameters of a DPP kernel is difficult due to the non-convexity of the likelihood. We proposed Bayesian approaches using MCMC, in particular, for inferring these parameters. In addition to being more robust and providing a characterization of the posterior uncertainty, these algorithms can be modified to deal with large-scale and continuous DPPs. We also showed how our posterior samples can be evaluated using mo-

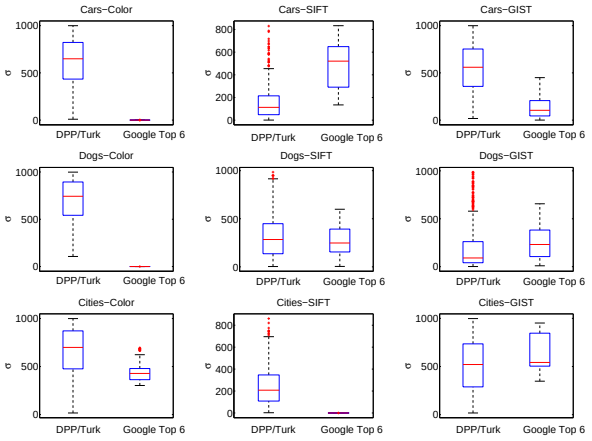


*Figure 6.* For the image diversity experiment, boxplots of posterior samples of (rom left to right) $\sigma^{\mathsf{cat}}_{\mathsf{color}}$, $\sigma^{\mathsf{cat}}_{\mathsf{SIFT}}$ and $\sigma^{\mathsf{cat}}_{\mathsf{GIST}}$. Each plot shows results for human annotated sets (left) versus Google Top 6 (right). Categories from top to bottom: (a) cars, (b) dogs and (c) cities.

ment matching as a model-checking method. Finally we demonstrated the utility of learning DPP parameters in studying diabetic neuropathy and evaluating human perception of diversity in images.

## References


Affandi, R. H., Kulesza, A., and Fox, E. B. Markov determinantal point processes. In *Proc. UAI*, 2012.

Affandi, R. H., Fox, E.B., and Taskar, B. Approximate inference in continuous determinantal processes. In *Proc. NIPS*, 2013a.

Affandi, R.H., Kulesza, A., Fox, E.B., and Taskar, B. Nyström approximation for large-scale determinantal processes. In *Proc. AISTATS*, 2013b.

Borodin, A. Determinantal point processes. *arXiv preprint arXiv:0911.1153*, 2009.

Borodin, A. and Rains, E.M. Eynard-Mehta theorem, Schur process, and their Pfaffian analogs. *Journal of statistical physics*, 121(3):291–317, 2005.

Fasshauer, G.E. and McCourt, M.J. Stable evaluation of Gaussian radial basis function interpolants. *SIAM Journal on Scientific Computing*, 34(2):737–762, 2012.

Gelman, A. and Rubin, D. B. Inference from iterative simulation using multiple sequences. *Statistical science*, pp. 457–472, 1992.

Gillenwater, J., Kulesza, A., and Taskar, B. Discovering diverse and salient threads in document collections. In *Proc. EMNLP*, 2012.

Guan, K. Schur-convexity of the complete elementary symmetric function. *Journal of Inequalities and Applications*, 2006(1):67624, 2006.

Hough, J.B., Krishnapur, M., Peres, Y., and Virág, B. Determinantal processes and independence. *Probability Surveys*, 3:206–229, 2006.

Kulesza, A. and Taskar, B. Structured determinantal point processes. In *Proc. NIPS*, 2010.

Kulesza, A. and Taskar, B. k-DPPs: Fixed-size determinantal point processes. In *ICML*, 2011a.

Kulesza, A. and Taskar, B. Learning determinantal point processes. In *In Proc. UAI*, 2011b.

Kulesza, A. and Taskar, B. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*, 5(2–3), 2012.

Lavancier, F., Møller, J., and Rubak, E. Statistical aspects of determinantal point processes. *arXiv preprint arXiv:1205.4818*, 2012.

Lowe, D. G. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pp. 1150–1157. Ieee, 1999.

McKinnon, K. I.M. Convergence of the Nelder–Mead simplex method to a nonstationary point. *SIAM Journal on Optimization*, 9(1):148–158, 1998.

Neal, R. M. Slice sampling. *Annals of statistics*, pp. 705–741, 2003.

Nelder, J. and Mead, R. A simplex method for function minimization. *Computer Journal*, 7(4):308–313, 1965.

Oliva, A. and Torralba, A. Building the gist of a scene: The role of global image features in recognition. *Progress in brain research*, 155:23–36, 2006.

Paris, R.B. Asymptotics of integrals of hermite polynomials. *Appl. Math. Sci. 4*, pp. 3043–3056, 2010.

Ripley, B. D. Modelling spatial patterns. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 172–212, 1977.

Snoek, J., Zemel, R., and Adams, R. P. A determinantal point process latent variable model for inhibition in neural spiking data. In *Proc. NIPS*, 2013.

Vedaldi, A. and Fulkerson, B. Vlfeat: An open and portable library of computer vision algorithms. In *Proceedings of the international conference on Multimedia*, pp. 1469–1472. ACM, 2010.

Waller, L. A., Särkkä, A., Olsbo, V., Myllymäki, M., Panoutsopoulou, I.G., Kennedy, W.R., and Wendelschafer-Crabb, G. Second-order spatial analysis of epidermal nerve fibers. *Statistics in Medicine*, 30(23):2827–2841, 2011.

Zou, J. and Adams, R.P. Priors for diversity in generative latent variable models. In *Proc. NIPS*, 2012.

# Supplementary Material: Learning the Parameters of Determinantal Point Process Kernels

**Raja Hafiz Affandi** RAJARA@WHARTON.UPENN.EDU
University of Pennsylvania

**Emily B. Fox** EBFOX@STAT.WASHINGTON.EDU
University of Washington

**Ryan P. Adams** RPA@SEAS.HARVARD.EDU
Harvard University

**Ben Taskar** TASKAR@CS.WASHINGTON.EDU
University of Washington

## A. Gradient for Discrete DPP

Gradient ascent and stochastic gradient ascent provide attractive approaches in learning parameters, $\Theta$ of DPP kernel $L(\Theta)$ because of their theoretical guarantees, but require knowledge of the gradient of the log-likelihood $\mathcal{L}(\Theta)$. In the discrete DPP setting, this gradient can be computed straightforwardly and we provide examples for discrete Gaussian and polynomial kernels here.

$$\mathcal{L}(\Theta) = \sum_{t=1}^{T} \log(\det(L_{A^t}(\Theta))) - T\log(\det(L(\Theta)+I)) \tag{30}$$

$$\begin{aligned}\frac{d\mathcal{L}(\Theta)}{d\Theta} = &\sum_{t=1}^{T} \mathrm{tr}\left(L_{A^t}(\Theta)^{-1}\frac{dL_{A^t}(\Theta)}{d\Theta}\right)\\ &-T\mathrm{tr}\left((L(\Theta)+I)^{-1}\frac{dL(\Theta)}{d\Theta}\right)\end{aligned}$$

To find the MLE, we can perform gradient ascent

$$\Theta_i = \Theta_{i-1} + \eta\frac{dl(\Theta)}{d\Theta} \tag{31}$$

In the following examples, we denote $\mathbf{x}_i = (x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(d)})$, where $d$ is the number of dimension.

### A.1. Example I: Gaussian Similarity with Uniform Quality

$L(\Sigma) = \exp\{-(\mathbf{x}-\mathbf{y})^\top\Sigma^{-1}(\mathbf{x}-\mathbf{y})\}$

Denote $G_{ij}^{(lm)} = L_{ij}\frac{(x_i^{(l)}-x_j^{(l)})(x_i^{(m)}-x_j^{(m)})}{2\Sigma_{lm}^2}$

Then,

$$\begin{aligned}\frac{d\mathcal{L}(\Sigma)}{d\Sigma_{lm}} = &\sum_{t=1}^{T} \mathrm{tr}\left(L_{A^t}(\Sigma)^{-1}G_{A^t}^{(lm)}\right)\\ &-T\mathrm{tr}\left((L(\Sigma)+I)^{-1}G^{(lm)}\right)\end{aligned}$$

### A.2. Example II: Gaussian Similarity with Gaussian Quality

$L(\Gamma,\Sigma)$
$= \exp\{-\mathbf{x}^\top\Gamma^{-1}\mathbf{x} - (\mathbf{x}-\mathbf{y})^\top\Sigma^{-1}(\mathbf{x}-\mathbf{y}) - \mathbf{y}^\top\Gamma^{-1}\mathbf{y}\}$
Denote $C_{ij}^{(lm)} = L_{ij}\frac{(x_i^{(l)}x_i^{(m)}+x_j^{(l)}x_j^{(m)})}{2\Gamma_{lm}^2}$ and $G_{ij}^{(lm)}$ as in previous example.

Then,

$$\begin{aligned}\frac{d\mathcal{L}(\Gamma,\Sigma)}{d\Gamma_{lm}} = &\sum_{t=1}^{T} \mathrm{tr}\left(L_{A^t}(\Sigma)^{-1}C_{A^t}^{(lm)}\right)\\ &-T\mathrm{tr}\left((L(\Sigma)+I)^{-1}C^{(lm)}\right)\end{aligned}$$

and $\frac{dl(\Gamma,\Sigma)}{d\Sigma_{lm}}$ is the same as the previous example.

### A.3. Example III: Polynomial Similarity with Uniform Quality

$L(p,q) = \left(\mathbf{x}^\top\mathbf{y}+p\right)^q$

Denote $R_{ij} = qL_{ij}^{\frac{q-1}{q}}$ and $U_{ij} = L_{ij}\log(L_{ij}^{\frac{1}{q}})$.

Then,

$$\begin{aligned}\frac{d\mathcal{L}(p,q)}{dp} = &\sum_{t=1}^{T} \mathrm{tr}\left(L_{A^t}(p,q)^{-1}R_{A^t}\right)\\ &-T\mathrm{tr}\left((L(p,q)+I)^{-1}R\right)\end{aligned}$$

**Algorithm 1** Random-Walk Metropolis-Hastings

**Input**: Dimension: $D$, Starting point: $\Theta_0$, Prior distribution: $\mathcal{P}(\Theta)$, Proposal distribution $f(\hat{\Theta}|\Theta)$ with mean $\Theta$, Samples: $A^1, \ldots, A^T]$.
$\Theta = \Theta_0$
**for** $i = 0 : (\tau - 1)$ **do**
  $\hat{\Theta} \sim f(\hat{\Theta}|\Theta_i)$
  $r = \left(\frac{\mathcal{P}(\hat{\Theta}|A^1,\ldots,A^T)}{\mathcal{P}(\Theta_i|A^1,\ldots,A^T)} \frac{f(\Theta_i|\hat{\Theta})}{f(\hat{\Theta}|\Theta_i)}\right)$
  $u \sim$ Uniform[0,1]
  **if** $u < \min\{1, r\}$ **then**
    $\Theta_{i+1} = \hat{\Theta}$
**Output**: $\Theta_{0:\tau}$

$$\frac{d\mathcal{L}(p,q)}{dq} = \sum_{t=1}^{T} \mathrm{tr}\left(L_{A^t}(p,q)^{-1} U_{A^t}\right) - T\mathrm{tr}\left((L(p,q)+I)^{-1}U\right)$$

## B. Bayesian Learning

In the main paper, we highlight two techniques: random-walk Metropolis-Hastings (MH) and slice sampling to sample from the posterior distribution. We present the pseudo algorithms here (Alg. 1 and Alg. 2).

We also present the pseudo random walk MH algorithm for handling large-scale and continuous DPPs using posterior bounds in Alg. 3. We also present and illustration of the slice sampling using posterior bounds in Figure 7.

## C. Proof of DPP/$k$DPP Bounds

**Proposition C.1** *Let $\lambda_{1:\infty}$ be the eigenvalues of kernel $L$. Then*

$$\prod_{n=1}^{M}(1+\lambda_n) \leq \prod_{n=1}^{\infty}(1+\lambda_n) \tag{32}$$

*and*

$$\prod_{n=1}^{\infty}(1+\lambda_n) \leq \exp\left\{\mathrm{tr}(L) - \sum_{n=1}^{M}\lambda_n\right\}\left[\prod_{n=1}^{M}(1+\lambda_n)\right]. \tag{33}$$

**Proof:** The first inequality is trivial since the eigenvalues $\lambda_{1:\infty}$ are all nonnegative.

To proof the second inequality, we use the AM-GM inequality: For any non-negative numbers, $\gamma_1, \ldots, \gamma_M$, $(\prod_{n=1}^{M}\gamma_n)^{\frac{1}{M}} \leq \sum_{n=1}^{M}\frac{\gamma_n}{M}$.

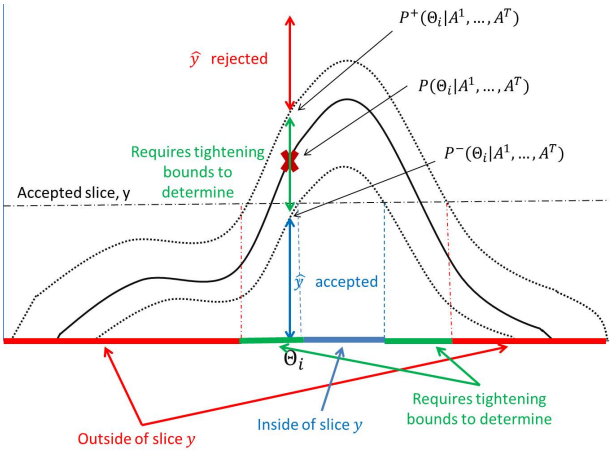


*Figure 7.* Illustration of slice sampling algorithm using posterior bounds. In the first step, a candidate slice $\hat{y}$ is generated. $\hat{y}$ is rejected if it is above the upper posterior bound and rejected if it is below the lower posterior bound. If $\hat{y}$ is in between the bounds, then the bounds are tightened until a decision can be made. Once a slice, $y$ is accepted, we need to sample new parameters inside the slice. To determine whether the endpoints of the interval or the new parameters are in the slice we decide that they are in the slice if the upper bound of posterior probability evaluated at the points are higher than the slice value and decide that they are outside of the slice if the lower bound of the posterior probability is lower than the slice value. Otherwise, we tighten the bounds until a decision can be made.

**Algorithm 2** Univariate Slice Sampling

**Input:** Starting point: $\Theta_0$, Initial width: $w$, Prior distribution: $\mathcal{P}(\Theta)$, Samples: $X = [X^1, \ldots, X^T]$.
$\Theta = \Theta_0$
**for** $i = 0 : (\tau - 1)$ **do**
  $y \sim \text{Uniform}[0, \mathcal{P}(\Theta_i | A^1, \ldots, A^T)]$
  $z \sim \text{Uniform}[0, 1]$
  $L = \Theta_i - z * \frac{w}{2}$
  $R = L + \frac{w}{2}$
  **while** $y > \mathcal{P}(L | A^1, \ldots, A^T)$ **do**
    $L = L - \frac{w}{2}$
  **while** $y > \mathcal{P}(R | A^1, \ldots, A^T)$ **do**
    $R = R + \frac{w}{2}$
  $\hat{\Theta} \sim \text{Uniform}[L, R]$
  **if** $\mathcal{P}(\hat{\Theta} | A^1, \ldots, A^T) < y$ **then**
    **while** $\mathcal{P}(\hat{\Theta} | A^1, \ldots, A^T) < y$ **do**
      **if** $\hat{\Theta} > \Theta$ **then**
        $R = \hat{\Theta}$
      **else**
        $L = \hat{\Theta}$
      $\hat{\Theta} \sim \text{Uniform}[L, R]$
  $\Theta_{i+1} = \hat{\Theta}$
**Output**: $\Theta_{0:\tau}$

Let $\Lambda_M = \sum_{n=M+1}^{\infty} \lambda_n$ and $\gamma_n = 1 + \lambda_n$. Then,

$$\begin{aligned}
\prod_{n=1}^{\infty}(1+\lambda_n) &= \prod_{n=1}^{\infty}\gamma_n = (\prod_{n=1}^{M}\gamma_n)(\prod_{n=M+1}^{\infty}\gamma_n) \\
&= (\prod_{n=1}^{M}\gamma_n)(\lim_{l\to\infty}\prod_{n=M+1}^{M+l}\gamma_n) \\
&\leq (\prod_{n=1}^{M}\gamma_n)(\lim_{l\to\infty}(\sum_{n=M+1}^{M+l}\frac{\gamma_n}{l})^l) \\
&\leq (\prod_{n=1}^{M}(1+\lambda_n))\exp(\Lambda_M) \ .
\end{aligned}$$

□

**Proposition C.2** *Let $\lambda_{1:\infty}$ be the eigenvalues of kernel $L$. Then*

$$e_k(\lambda_{1:M}) \leq e_k(\lambda_{1:\infty}) \tag{34}$$

*and*

$$e_k(\lambda_{1:\infty}) \leq \sum_{j=0}^{k} \frac{(\text{tr}(L) - \sum_{n=1}^{M}\lambda_n)^j}{j!} e_{k-j}(\lambda_{1:M}) \ . \tag{35}$$

**Proof:** Let $e_k(\lambda_{1:m})$ be the $k$th elementary symmetric function: $e_k(\lambda_{1:m}) = \sum_{J\subseteq\{1,\ldots,m\},|J|=k}\prod_{j\in J}\lambda_j$.

**Algorithm 3** Random-Walk Metropolis-Hastings with Posterior Bounds

Input: Dimension: $D$, , Starting point: $\Theta_0$, Prior distribution: $\mathcal{P}(\Theta)$, Proposal distribution $f(\hat{\Theta}|\Theta)$ with mean $\Theta$, samples: $X = [X^1, \ldots, X^T]$.
$\Theta = \Theta_0$
**for** $i = 0 : \tau$ **do**
  $\hat{\Theta} \sim f(\hat{\Theta}|\Theta_i)$
  $r_+ = \infty, r_- = -\infty$
  $u \sim \text{Uniform}[0,1]$
  **while** $u \in [r_-, r_+]$ **do**
    $r^+ = \left(\frac{\mathcal{P}^+(\hat{\Theta}|A^1,\ldots,A^T)}{\mathcal{P}^-(\Theta_i|A^1,\ldots,A^T)}\frac{f(\Theta_i|\hat{\Theta})}{f(\hat{\Theta}|\Theta_i)}\right)$
    $r^- = \left(\frac{\mathcal{P}^-(\hat{\Theta}|A^1,\ldots,A^T)}{\mathcal{P}^+(\Theta_i|A^1,\ldots,A^T)}\frac{f(\Theta_i|\hat{\Theta})}{f(\hat{\Theta}|\Theta_i)}\right)$
    Increase tightness on $\mathcal{P}^+$ and $\mathcal{P}^-$
  **if** $u < \min\{1, r^-\}$ **then**
    $\Theta_t = \hat{\Theta}$
Output: $\Theta_{0:\tau}$

Trivially, we have a lower bound since the eigenvalues $\lambda_{1:\infty}$ are non-negative: $e_k(\lambda_{1:m}) \leq e_k(\lambda_{1:n}) \quad \text{for } m \leq n$.

For the upper bound we can use the Schur-concavity of elementary symmetric functions for non-negative arguments (Guan, 2006).Thus for $\bar{\lambda}_{1:N} \prec \lambda_{1:N}$:

$$\sum_{i=1}^{k}\bar{\lambda}_n \leq \sum_{n=1}^{k}\lambda_n \quad \text{for } k = 1, \ldots, N-1 \tag{36}$$

and

$$\sum_{n=1}^{N}\bar{\lambda}_n = \sum_{n=1}^{N}\lambda_n \tag{37}$$

we have $e_k(\bar{\lambda}_{1:N}) \geq e_k(\lambda_{1:N})$.

Now let $\Lambda_M = \sum_{n=M+1}^{\infty}$ and $\Lambda_M^N = \sum_{n=M+1}^{N}$. We consider

$$\bar{\lambda}_{1:N}^{(M)} = (\lambda_1, \ldots, \lambda_M, \frac{\Lambda_M^N}{N-M}, \ldots, \frac{\Lambda_M^N}{N-M}). \tag{38}$$

Note that $\bar{\lambda}_{1:N}^{(M)} \prec \lambda_{1:N}$ and so $e_k(\bar{\lambda}_{1:N}^{(M)}) \geq e_k(\lambda_{1:N})$ for $M < N$.

We now compute $e_k(\bar{\lambda}_{1:N}^{(M)})$. Note that for $e_k(\bar{\lambda}_{1:N}^{(M)})$, the terms in the sum are products of $k$ factors, each containing some of the $\lambda_{1:M}$ factors and some of the $\frac{\Lambda_M^N}{N-M}$ factors. The sum of the terms that have $j$ factors of type $\frac{\Lambda_M^N}{N-M}$ is $\binom{N-M}{j}\left(\frac{\Lambda_M^N}{N-M}\right)^j e_{k-j}(\Lambda(m))$, so we have:

$$e_k(\bar{\lambda}_{1:N}^{(M)}) = \sum_{j=0}^{k}\binom{N-M}{j}\left(\frac{\Lambda_M^N}{N-M}\right)^j e_{k-j}(\lambda_{1:M}) \ .$$

Using $\binom{N-M}{j} \leq \frac{(N-M)^j}{j!}$, we get

$$e_k(\bar{\lambda}^{(M)}_{1:N}) = \sum_{j=0}^{k} \left( \frac{(\Lambda_M^N)^j}{j!} \right) e_{k-j}(\lambda_{1:M}) \ .$$

Letting $N \to \infty$, we get out upper bound

$$e_k(\lambda_{1:\infty}) \leq \sum_{j=0}^{k} \left( \frac{(\Lambda_M)^j}{j!} \right) e_{k-j}(\lambda_{1:M}) \quad \text{for } m \leq n.$$

## D. Moments for Continuous DPP with Gaussian Quality and Similarity

In the continuous case, given the eigendecomposition of the kernel operator, $L(\mathbf{x}, \mathbf{y}) = \sum_{n=1}^{\infty} \lambda_n \phi_n(\mathbf{x})^* \phi_n(\mathbf{y})$ (where $\phi_n(\mathbf{x})^*$ denotes the complex conjugate of the $n$th eigenfunction), the $m$th moment can be evaluated as

$$E[\mathbf{x}^m] = \int_{\Omega} \sum_{n=1}^{\infty} \frac{\lambda_n}{\lambda_n + 1} \mathbf{x}^m \phi_n(\mathbf{x})^2 d\mathbf{x} \ . \tag{39}$$

Here we present the derivation for moments when

$$q(\mathbf{x}) = \sqrt{\alpha} \prod_{d=1}^{D} \frac{1}{\sqrt{\pi \rho_d}} \exp\left\{ -\frac{x_d^2}{2\rho_d} \right\} \tag{40}$$

and

$$k(\mathbf{x}, \mathbf{y}) = \prod_{d=1}^{D} \exp\left\{ -\frac{(x_d - y_d)^2}{2\sigma_d} \right\}, \mathbf{x}, \mathbf{y} \in \mathbb{R}^D. \tag{41}$$

In this case, the eigenvalues and eigenvectors of the operator $L$ are given by Fasshauer & McCourt (2012),

$$\lambda_{\mathbf{n}} = \alpha \prod_{d=1}^{D} \sqrt{\frac{1}{\frac{\beta_d^2+1}{2} + \frac{1}{2\gamma_d}}} \left( \frac{1}{\gamma_d(\beta_d^2+1)+1} \right)^{n_d - 1}, \tag{42}$$

and

$$\phi_{\mathbf{n}}(\mathbf{x}) = \prod_{d=1}^{D} \left( \frac{1}{\pi \rho_d^2} \right)^{\frac{1}{4}} \sqrt{\frac{\beta_d}{2^{n_d-1}\Gamma(n_d)}} \exp\left\{ -\frac{\beta_d^2 x^2}{2\rho_d^2} \right\} H_{n_d - 1}\left( \frac{\beta_d x_d}{\sqrt{\rho_d^2}} \right) \tag{43}$$

where $\gamma_d = \frac{\sigma_d}{\rho_d}$ , $\beta_d = (1 + \frac{2}{\gamma_d})^{\frac{1}{4}}$ and $\mathbf{n} = (n_1, n_2, \ldots, n_D)$ is a multi index.

In the case of DPPs (as opposed to $k$-DPPs), we can use the number of items as an estimate of the 0th moment. The 0th moment is given by $\sum_{\mathbf{n}=1} \frac{\lambda_{\mathbf{n}}}{1+\lambda_{\mathbf{n}}}$.

Denote $\mathbf{x} = (x_1, x_2, \ldots, x_d)$. The for higher moments, note that

$$\begin{aligned} E[x_j^m] &= \int_{\mathbb{R}} \sum_{n=1}^{\infty} \frac{\lambda_n}{\lambda_n + 1} x_j^m \phi_n(\mathbf{x})^2 dx_j \\ &= \sum_{n=1}^{\infty} \frac{\lambda_n}{\lambda_n + 1} \int_{\mathbb{R}} x_j^m \phi_n(\mathbf{x})^2 dx_j \end{aligned}$$

Using the results of moment integrals involving a product of two Hermite polynomials (Paris, 2010), we get that

$$E[x_j^m] = \int_{\mathbb{R}^d} \sum_{\mathbf{n}}^{\infty} \frac{\lambda_{\mathbf{n}}}{\lambda_{\mathbf{n}} + 1} \left( \frac{\sqrt{\rho_j}}{\sqrt{2}\beta_j} \right)^m \wp_{\frac{m}{2}}(n_j - 1) \tag{44}$$

for $m$ even and 0 otherwise. The polynomial $\wp_{\frac{m}{2}}(n_j - 1)$ is given in Eq. (4.8) in Paris (2010). For example, the second and fourth moments are given by

(i) $E[x_j^2] = \sum_{\mathbf{n}}^{\infty} \frac{\lambda_{\mathbf{n}}}{\lambda_{\mathbf{n}}+1} \left( \frac{\sqrt{\rho_j}}{\sqrt{2}\beta_j} \right)^2 (2n_j - 1)$

(ii) $E[x_j^4] = \sum_{\mathbf{n}}^{\infty} \frac{\lambda_{\mathbf{n}}}{\lambda_{\mathbf{n}}+1} \left( \frac{\sqrt{\rho_j}}{\sqrt{2}\beta_j} \right)^4 3(2n_j^2 - 2n_j + 1)$

For a low dimensional setting, we can learn the parameters by using grid search such that the moments agree.

## E. Details on Simulation

In the main paper, we use our Bayesian learning algorithms to learn parameters from (i) simulated data generated from a 2-dimensional isotropic discrete kernel ($\sigma_d = \sigma$, $\rho_d = \rho$ for $d = 1, 2$), (ii) nerve fiber data using 2-dimensional isotropic continuous kernel ($\sigma_d = \sigma$, $\rho_d = \rho$ for $d = 1, 2$) and (ii) image diversity data using 3600-dimensional discrete kernel with Gaussian similarity. In all of these experiments, we use weakly informative inverse gamma priors on $\sigma$,$\rho$ and $\alpha$. In particular, for all three parameters, we used the same priors for all three parameters

$$\mathcal{P}(\alpha) = \mathcal{P}(\rho) = \mathcal{P}(\sigma) = \text{Inv-Gamma}(0.001, 0.001) \tag{45}$$

We then learn the parameters using hyperrectangle slice sampling.

## F. Details on Image Diversity

In studying the diversity in images, we extracted 3 different types of features from the images—color features, SIFT-descriptors (Lowe, 1999; Vedaldi & Fulkerson, 2010) and GIST-descriptors (Oliva & Torralba, 2006) described in the supplementary material. We describe these features below.

**Color**: Each pixel is assigned a coordinate in three-dimensional Lab color space. The colors are then sorted

into axis-aligned bins, producing a histogram of either 8 (denoted color8) or 64 (denoted color64) dimensions.

**SIFT**: The images are processed to obtain sets of 128-dimensional SIFT descriptors. These descriptors are commonly used in object recognition to identify objects in images and are invariant to scaling, orientation and minor distortions. The descriptors for a given category are combined, subsampled to set of 25,000, and then clustered using k-means into either 256 (denoted SIFT256) or 512 (denoted SIFT512) clusters. The feature vector for an image is the normalized histogram of the nearest clusters to the descriptors in the image.

**GIST**: The images are processed to obtain 960-dimensional GIST feature vectors that is commonly used to describe scene structure.

We also extracted the features above from the center of the images, defined as the centered rectangle with dimensions half those of the original image. This yields a total of 10 different feature vectors. Since we are only concerned with the diversity of the images, we ensure that the quality across the images are uniform by normalizing each feature vector such that their $L_2$ norm equals to 1. We then combine the feature vectors into 3 types of features- color, SIFT and GIST.

For the Google top 6 images, we model the samples, $X^t_{Top6}$ as though they are generated from a 6-DPP with kernel $L^{subcat}(X^t)$. To highlight the effect of the human annotation in the partial results sets, we model the samples as though they are generated from a conditional 6-DPP.

In general, given a partial set of observations A and $k$-DPP kernel $L$, we can define the conditional $k$-DPP probability of choosing a set B given the inclusion of set A (with $|A| + |B| = k$)as

$$\mathcal{P}^k_L(Y = A \cup B | A \in Y) \propto \det(L^A_B) \tag{46}$$

with

$$L^A = \left( \left[ ((L + I_{A^c})^{-1} \right]_{A^c} \right)^{-1} - I \tag{47}$$

where $I_{A^c}$ denotes the identity matrix with 0 for diagonal corresponding to elements in $A$. Here, following the $N \times N$ inversion, the matrix is restricted to rows and columns indexed by elements not in $A$, then inverted again. The normalizer is given by Kulesza & Taskar (2012).

$$\sum_{|Y'|=k-|A|} \det(L^A_{Y'}) \tag{48}$$

In our experiment, our samples can be seperated into the partial result sets and human annotations,

$$X^t_{\text{DPP+human}} = (A^t, b^t) \tag{49}$$

where $A^t$ is the partial result sets and $b^t$ is the human annotated result, we model the data from the conditional 6-DPP $L^{subcat}(b^t|A^t)$. In this case, the likelihood is given by

$$L^i(\Theta^{cat}) = \frac{\det(L^{i\ A^t}_{b^t_t}(\Theta^{cat}))}{\sum_{i=1}^N L^{i\ A^t}_{x_i}(\Theta^{cat})} \tag{50}$$

for each subcategory, $i$. That is, for each subcategory, i, we compute $L^i(\Theta^{cat})$ and use Eq. 47 to compute the conditional kernel.